Zabbix Manual

Welcome to the user manual for Zabbix 2.0 software. These pages are created to help our users successfully manage their monitoring tasks with Zabbix, from the simple to the more complex.

2.0/manual.txt · Last modified: 2012/05/23 11:27 by martins-v

Except where otherwise noted, content on this wiki is licensed under the following license:CC Attribution-Noncommercial-Share Alike 3.0 Unported [http://creativecommons.org/licenses/by-nc-sa/3.0/]

14. Distributed monitoring

Overview

Zabbix provides effective and reliable ways of monitoring distributed IT infrastructure. The two main solutions for large environments provided by Zabbix are:

- use of proxies
- use of nodes

Proxies can be used to collect data locally on behalf of a centralized Zabbix server and then report the data to the server. Nodes are full Zabbix servers that can be set up in a hierarchy of distributed monitoring.

Proxy vs. node

When making a choice between using a proxy or a node, several considerations must be taken into account.

	Proxy	**Node**
Lightweight	**Yes**	No
GUI	No	**Yes**
Works independently	**Yes**	**Yes**
Easy maintenance	**Yes**	No
Automatic DB creation[1]	**Yes**	No
Local administration	No	**Yes**
Ready for embedded hardware	**Yes**	No
One way TCP connections	**Yes**	**Yes**
Centralised configuration	**Yes**	No
Generates notifications	No	**Yes**

[1] Automatic DB creation feature only works with SQLite. Other databases require a manual setup.

2.0/manual/distributed_monitoring.txt · Last modified: 2011/12/27 17:55 by martins-v
Except where otherwise noted, content on this wiki is licensed under the following license:CC Attribution-Noncommercial-Share Alike 3.0 Unported [http://creativecommons.org/licenses/by-nc-sa/3.0/]

1 Proxies

Overview

A Zabbix proxy can collect performance and availability data on behalf of the Zabbix server. This way, a proxy can take on itself some of the load of collecting data and offload the Zabbix server.

Also, using a proxy is the easiest way of implementing centralized and distributed monitoring, when all agents and proxies report to one Zabbix server and all data is collected centrally.

A Zabbix proxy can be used to:

- Monitor remote locations
- Monitor locations having unreliable communications
- Offload the Zabbix server when monitoring thousands of devices
- Simplify the maintenance of distributed monitoring

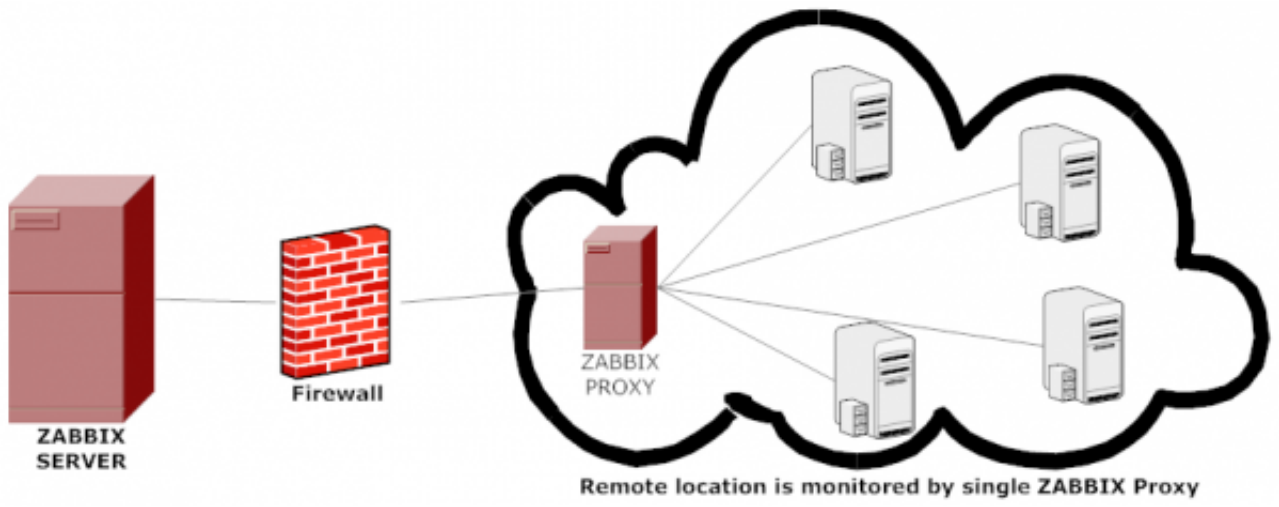

The proxy requires only one TCP connection to the Zabbix server. This way it is easier to get around a firewall as you only need to configure one firewall rule.

Zabbix proxy must use a separate database. Pointing it to the Zabbix server database will break the configuration.

All data collected by the proxy is stored locally before transmitting it over to the server. This way no data is lost due to any temporary communication problems with the server. The *ProxyLocalBuffer* and *ProxyOfflineBuffer* parameters in the proxy configuration file control for how long the data are kept locally.

Zabbix proxy is a data collector. It does not calculate triggers, process events or send alerts. For an overview of what proxy functionality is, review the following table:

Function	Supported by proxy
Items	
Zabbix agent checks	**Yes**
Zabbix agent checks (active)	**Yes** [1]

Simple checks	**Yes**
Trapper items	**Yes**
SNMP checks	**Yes**
SNMP traps	**Yes**
IPMI checks	**Yes**
JMX checks	**Yes**
Log file monitoring	**Yes**
Internal checks	*No*
SSH checks	**Yes**
Telnet checks	**Yes**
External checks	**Yes**
Built-in web monitoring	**Yes**
Network discovery	**Yes**
Low-level discovery	**Yes**
Calculating triggers	*No*
Processing events	*No*
Sending alerts	*No*
Remote commands	*No*

[1] To make sure that an agent asks the proxy (and not the server) for active checks, the proxy must be listed in the **ServerActive** parameter in the agent configuration file.

Configuration

Once you have installed and configured a proxy, it is time to configure it in the Zabbix frontend.

Adding proxies

To configure a proxy in Zabbix front end:

- Go to: *Administration → DM*
- Select *Proxies* from the dropdown in the upper right corner
- Click on *Create proxy* (or on the name of an existing proxy)

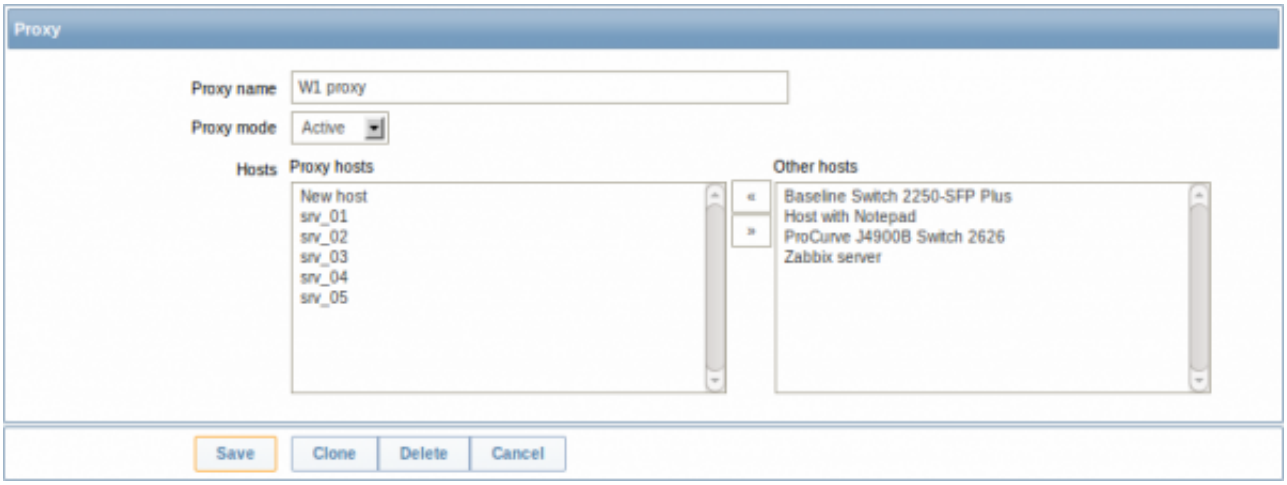

Parameter	Description
Proxy name	Enter the proxy name. It must be the same name as in the *Hostname* parameter in the proxy configuration file.
Proxy mode	Select the proxy mode. **Active** - the proxy will connect to the Zabbix server and request configuration data **Passive** - Zabbix server connects to the proxy
Hosts	Add hosts to be monitored by the proxy.

Host configuration

You can specify that an individual host should be monitored by a proxy in the host configuration form, using the *Monitored by proxy* field.

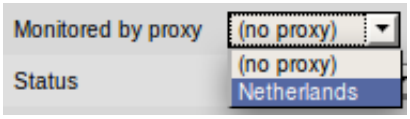

2.0/manual/distributed_monitoring/proxies.txt · Last modified: 2012/09/28 11:48 by martins-v

Except where otherwise noted, content on this wiki is licensed under the following license:CC Attribution-Noncommercial-Share Alike 3.0 Unported [http://creativecommons.org/licenses/by-nc-sa/3.0/]

2 Nodes

Overview

You can use nodes to build a hierarchy of distributed monitoring.

Each node is a full Zabbix server and is responsible for monitoring its own location. Zabbix supports up to a thousand nodes in a distributed setup.

The benefits of using a node setup:

- building a multi-level hierarchy of monitoring in a large network involving several geographical locations. A node in the hierarchy reports to its master node only.
- a node can be configured locally or through its master node, which has a copy of configuration data of all child nodes.
- data gathering becomes more immune to possible communication problems. If communication between a master and a child node breaks down, nodes can keep operating. Historical information and events are stored locally. When communication is back, a child node will optionally send the data to the master node.
- the nodes can split the work of a single Zabbix server having to monitor thousands of hosts
- attaching and detaching new nodes does not affect the functionality of the existing setup. No restart of any node required.

Platform independence

A node may use its own platform (OS, hardware) and database engine independently of other nodes. Also child nodes can be installed without Zabbix frontend.

The nodes of higher levels should use a combination of better hardware with MySQL InnoDB, Oracle or PostgreSQL backend.

A distributed monitoring setup will not work with an SQLite backend database.

Configuration

Node configuration

A Zabbix server installed by following the standard installation procedure is not configured as a node for a distributed setup.

To configure it as a node:

- Specify a unique **NodeID** in the server configuration file (zabbix_server.conf). Available values: 1-999 ('0' being the default value of a standalone server)

- Stop zabbix_server, make sure that is is NOT running
- Convert database data for a distributed setup, by running:

```
zabbix_server -n <node id>
```

Run this command only **once**. Running it twice will corrupt the database, so make sure that you run it with the correct node id.

For example, you may run (if NodeID is '1'):

```
cd bin
./zabbix_server -n 1 -c /usr/local/etc/zabbix_server.conf
```

Running zabbix_server with the **-n** option does not start the server process.

In a very simple setup, we may envisage this node (with NodeID=1) as the master, and go on to configure another Zabbix server as a child node, using the same procedure, only using a different node identifier, say, '2'. With two nodes configured, it is time to add them in the Zabbix front-end, in a very simple master-child relationship.

Front-end configuration (master node)

To configure the master node, open its Zabbix frontend:

- Go to: *Administration → DM*
- Make sure that *Nodes* are selected in the dropdown to the right
- Click on *Local node* to review its parameters

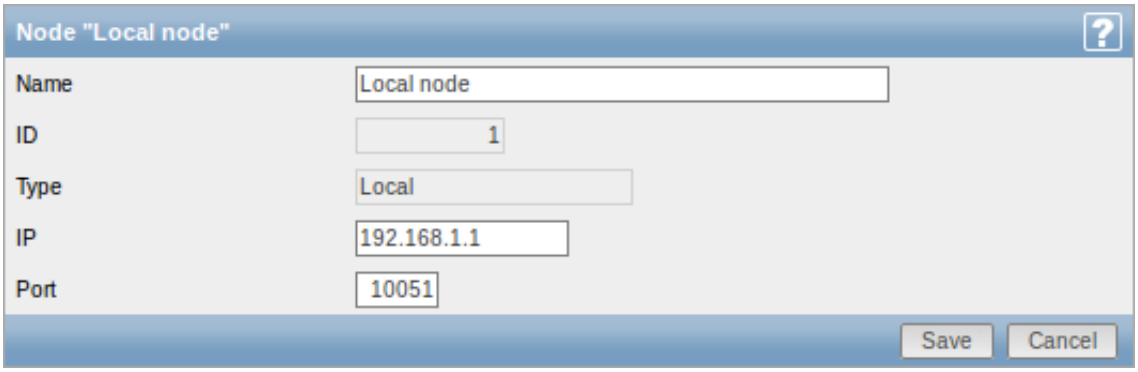

Node attributes:

Parameter	Description
Name	Unique node name.
Id	Unique node ID. This is the value of **NodeID** from the configuration file.
Type	**Local** - the local node
IP	IP address of the local node. Zabbix trapper must be listening on this IP address.
Port	Port number of the local node. Zabbix trapper must be listening on this port number. Default is 10051.

Then add the child node:

- Click on *New node* in *Administration → DM*

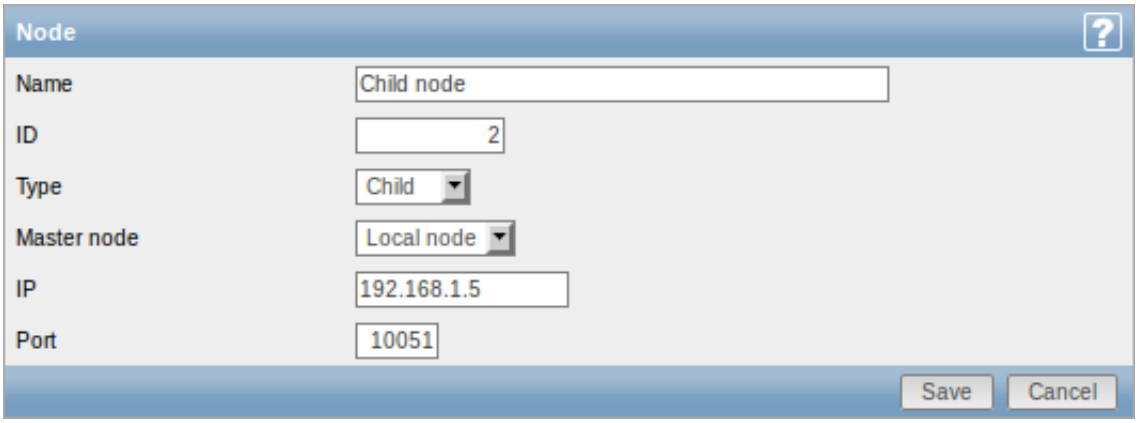

Node attributes:

Name	Unique name of the child node.
Id	Unique node ID. This is the **NodeID** from the child node configuration file.
Type	Select the first of the two available values: **Child** - a child node **Master** - a master node
Master node	Select the master node for this child node.
IP	IP address of the child node. Zabbix trapper must be listening on this IP address.
Port	Port number of the child node. Zabbix trapper must be listening on this port number. Default is 10051.

Front-end configuration (child node)

To configure the child node, open its Zabbix frontend:

- Go to: *Administration → DM*
- Make sure that *Nodes* are selected in the dropdown to the right
- Click on *Local node* to review its parameters (see above for how to configure the local node)

Then add the master node:

- Click on *New node* in *Administration → DM*

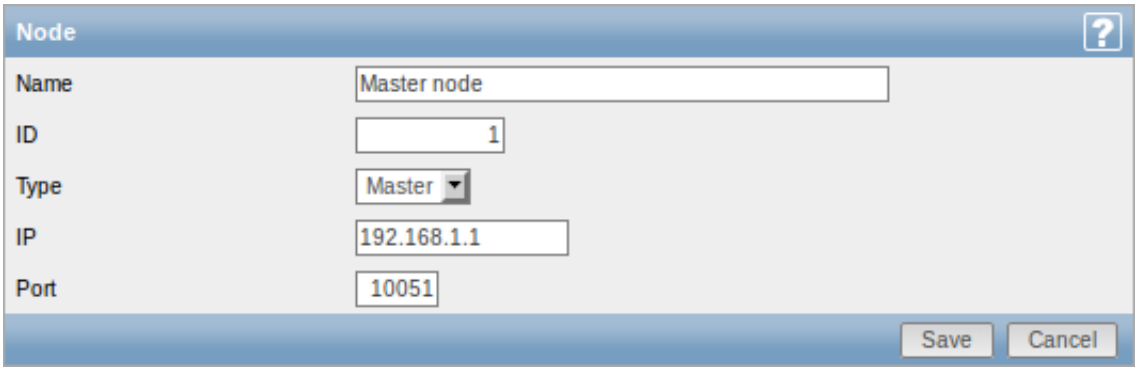

Node attributes:

Name	Unique name of the master node.

Id	Unique master node ID. This is the **NodeID** from the master node configuration file.
Type	Select the second of the two available values: **Child** - a child node **Master** - a master node
IP	IP address of the master node. Zabbix trapper must be listening on this IP address on the master node.
Port	Port number of the master node. Zabbix trapper must be listening on this port number. Default is 10051.

Starting server daemons

To finish configuring a simple distributed setup, start the zabbix_server daemons, beginning with the master node daemon.

Display

As soon as nodes are defined, a dropdown for selecting one or several nodes appears in the Zabbix front-end.

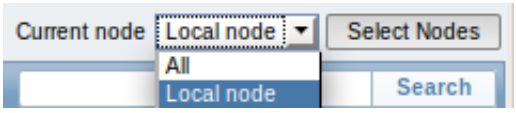

Once selected, all information displayed in the front-end will come from the selected node(s).

More complex configurations

You can use the principles outlined to build more complex, multi-level monitoring hierarchies.

In this example, Riga (Node 4) will collect events and history from all the child nodes.

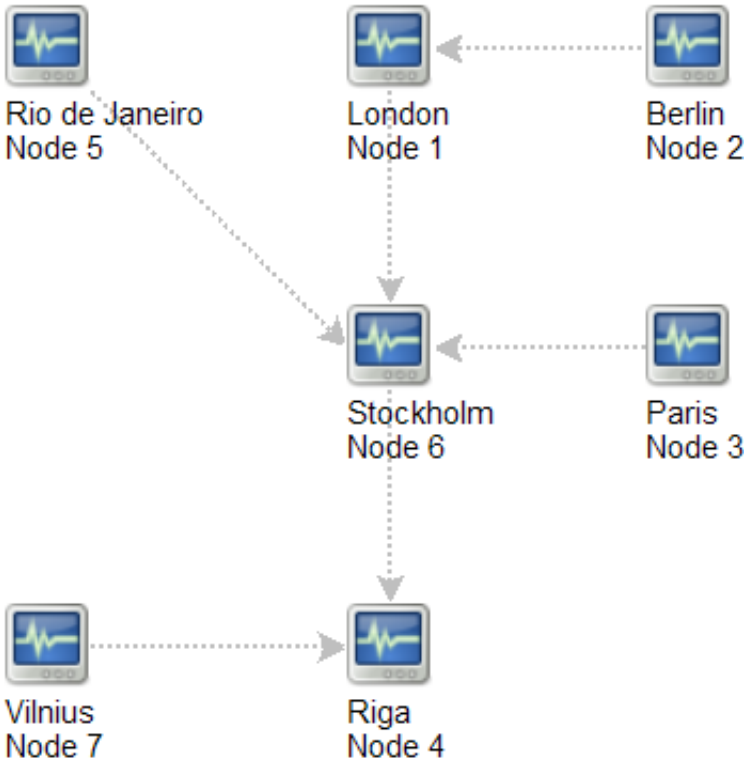

2.0/manual/distributed_monitoring/nodes.txt · Last modified: 2012/07/10 11:11 by martins-v

Except where otherwise noted, content on this wiki is licensed under the following license:CC Attribution-Noncommercial-Share Alike 3.0 Unported [http://creativecommons.org/licenses/by-nc-sa/3.0/]